

I hereby certify that I correspondence is being deposited
as express mail # EL102854092US
with the United States Postal Service in an envelope
addressed to: Commissioner of Patents and Trademarks,
Washington, D.C. 20231 on 6/25/99
Date

By [Signature]
6/25/99
Date of Signature

EL102854092US

Multi-Function High-Speed Network Interface

Field of the Invention

The current data communication invention is directed to a high speed network interface that is capable of the transmission and reception of multiple types of variable length data packets. In addition, the invention includes a media access control mechanism. Data packets can be sent over a variety of physical media, including single mode fiber, multi-mode fiber, and parallel fibers.

Background of the Invention

In data communications networks, a large number of methods are used to encapsulate communication data packets such as OSI (Open Systems Interconnect) layer 3 TCP/IP packets for the purpose of transmission over local or layer 2 networks and over specific point to point layer 1 physical links. Examples of OSI layer 2 network encapsulation and transmission and access control methods include 10Mb

1 Ethernet, 100Mb Ethernet, gigabit Ethernet (IEEE 802.3z),
2 IEEE 802.1Q, IEEE 802.3x, FDDI (ANSI X3T9.5), and token ring
3 (IEEE 802.5). There are also a large number of methods to
4 encapsulate layer 3 packets for transmission over layer 2
5 networks. For point to point links, available encapsulations
6 include PPP over HDLC (RFC1661), and Packet over Sonet (IETF
7 RFC1619).

8 In a generic Ethernet packet, the header information
9 that is immediately needed for a switching decision is the
10 layer 2 media access control (MAC) source and destination
11 addresses and, optional 802.1Q tag information. For an IP
12 packet, routing information is contained in the IP source
13 and destination addresses. The MAC source and destination
14 addresses are used for layer 2 switching, wherein the
15 destination address is matched with the port having
16 previously received a source address of the same value. The
17 layer 2 source and destination information is readily
18 available in the first 12 bytes of the Ethernet packet, and
19 generally presents no challenge in extraction. Higher level
20 layer 3 Internet Protocol (IP), and other types of protocol
21 packets present somewhat greater difficulty. For IP, a 32
22 bit IP source and 32 bit IP destination address are needed
23 for the routing decision, and the hardware must go through a
24 decision tree to determine what type of packet is being
25 examined, what protocol type it is, and thereafter extract

1 addressing information. Virtual Local Area Network (VLAN)
2 information may be added indicating which VLAN the packet
3 has membership. Each switch keeps a copy of a table with
4 the VLAN value associated with a port of exit. When a
5 packet bearing this VLAN value arrives, a hardware-based
6 lookup is performed into a table, which yields the
7 associated port of exit for the packet. If the decision
8 tree determines that the packet is not of a supported type,
9 then the treatment reverts back to the layer 2 switching
10 treatment of a simple Ethernet packet, and the time spent
11 examining the packet is lost. Additionally, for each
12 separate protocol, such as IP, IPX, Appletalk, etc., this
13 examination of the packet must occur, and the information
14 will appear in different places and formats in the packet.
15 It would be useful to have a single method of access for the
16 transmission and reception of switching and routing
17 information for such data packets including a common header
18 format.

19

20 Objects of the Invention

21 A first object of the invention is to provide a packet
22 encapsulation that includes a start symbol and a type field
23 which provides for the encapsulation of a variety of data
24 packet types, including Ethernet, FDDI, Token Ring, ATM
25 cells, and others.

1 A second object of the invention is a packet
2 encapsulation that provides for layer 2 prioritization and
3 virtual LAN information for any type of data packet.

4 A third object of the invention is the provision of a
5 header field which allows application specific packet
6 information in addition to the payload.

7 A fourth object of the invention is the provision of a
8 CRC that includes the header and payload of the packet.

9 A fifth object of the invention is to provide for data
10 transmission at a variety of speeds.

11 A sixth object of the invention is the provision of a
12 flow control mechanism to the media as indicated by the
13 receiver.

14 A seventh object of the invention is to allow data
15 transmission in a variety of encodings in a transmission and
16 reception media, including a serial channel, as well as any
17 number of parallel channels.

18

19 Summary of the Invention

20 The present invention comprises a method for encoding
21 and decoding data referred to as GX packets, a transmit
22 processor, and a receive processor. The transmit processor
23 includes a transmit buffer/controller dividing the data into
24 a plurality of transmit data lanes, a plurality of transmit
25 encoders, each encoder accepting information from a unique

1 transmit data lane and producing an output stream of clock-
2 encoded data, and a plurality of transmit serializers, each
3 accepting a unique stream of transmit encoded data, and
4 serially encoding it for transmission. A variable-speed
5 transmit clock circuit affords clocking of the elements of
6 the transmit processor. The transmit encoders add clock
7 recovery information to the data stream, optionally using
8 the 8B/10B encoding method, and also insert START, END,
9 IDLE_EVEN, IDLE_ODD, IDLE_EVEN_BUSY, and IDLE_ODD_BUSY
10 information as instructed by the input buffer/controller.
11 The receive processor comprises a plurality of receive
12 deserializers, each receiving a serial stream of encoded
13 data, and producing a stream of parallel encoded data.
14 These encoded parallel streams are passed on to a plurality
15 of receive decoders, each of which decodes a stream of
16 encoded data into a stream of byte information, as well as
17 decoding control information such as START, END, IDLE_EVEN,
18 IDLE_ODD, IDLE_EVEN_BUSY, and IDLE_ODD_BUSY. The receive
19 processor/controller accepts these streams of data and
20 control signals, organizes the recovered byte streams into
21 recovered packets, and also reports errors associated with
22 the data recovery process. A recovered receive clock for
23 each data lane is used to clock synchronized data into the
24 elasticity buffer of the receive processor.
25

1

2 Brief Description of the Drawings

3 Figure 1 is a prior art IEEE 802.3 Ethernet Packet
4 including an IP payload.

5 Figure 2 is a prior art ATM cell.

6 Figure 3a shows the GX packet format with header
7 details.

8 Figure 3b shows the header for the GX packet.

9 Figure 4 shows a data stream including packets and
10 inter-packet intervals.

11 Figure 5 shows the GX packet format with payload
12 details.

13 Figure 6 shows the case where the GX payload is an
14 Ethernet packet.

15 Figure 7 shows the case where the GX payload is a
16 native IP packet.

17 Figure 8 shows the case where the GX payload is an ATM
18 cell.

19 Figure 9 shows the case where the GX payload is an FDDI
20 packet.

21 Figure 10 shows the case where the GX payload is a
22 Token Ring packet.

23 Figure 11 shows the GX packet format where the data is
24 transmitted and received across 8 data lanes.

1 Figure 12 shows the GX packet format where the data is
2 transmitted and received across 4 data lanes.

3 Figure 13 shows the GX packet format where the data is
4 transmitted and received across 2 data lanes.

5 Figure 14 shows the GX packet format where the data is
6 transmitted and received across 1 data lane.

7 Figure 15 is an 8 data lane transmit processor for the
8 packet of figure 11.

9 Figure 16 is a 4 data lane transmit processor for the
10 packet of figure 12.

11 Figure 17a shows the 8B/10B encoder.

12 Figure 17b shows the encoding scheme for the encoder of
13 figure 17a.

14 Figure 18 is an 8 data lane receive processor for the
15 packet of figure 11.

16 Figure 19 is a 4 data lane receive processor for figure
17 12.

18 Figure 20a shows the 10B/8B decoder.

19 Figure 20b shows the decoding scheme for the decoder of
20 figure 20a.

21

22

23 Detailed Description of the Invention

24 Figure 1 shows a prior art layer 2 Ethernet Packet 10,
25 comprising a preamble 12, a header 16 comprising a MAC

1 Destination Address 14, a MAC Source Address 18, and
2 length/type information 22. Payload 24 is followed by the
3 Frame Check Sequence 26 which is a Cyclical Redundancy Check
4 (CRC) polynomial applied over the entire Ethernet header 16
5 and payload 24. For a generic layer 2 Ethernet packet,
6 payload 24 contains the data. In the case of a layer 3
7 protocol such as IP, payload 24 is arranged to further
8 comprise an IP header 28, an IP source address 30, and an IP
9 destination address 32, followed by the IP data 34. Other
10 layer 3 protocols such as Appletalk, IPX, and the like have
11 alternate arrangements for payload 24, but in general carry
12 a layer-related source and destination address observed by
13 the particular protocol. Other attributes of Ethernet
14 packet 10 include variable length payload 24, which may vary
15 from 46 bytes to 1500 bytes, as defined in the MAC layer
16 specification IEEE 802.3. The general attributes of prior
17 art packets as described in figure 1 include both layer 2
18 (MAC) and layer 3 (network) source and destination
19 addresses, and variable length data 24. These are used
20 individually, or in combination, by network switches and
21 routers to forward packets across layer 2 subnets, and
22 represents information for which the switching hardware
23 generally needs immediate access for the switching/routing
24 decision.

1 Figure 2 shows an ATM cell 40 generally used in an ATM
2 switching network. ATM networks set up end-to-end
3 connections according to ATM Forum User Network Interface
4 (UNI 3.1) protocols prior to the transmission of actual
5 network data. The characteristic of ATM cells is a fixed 53
6 byte cell comprising a 5 byte ATM header 31 and a 48 byte
7 payload 50. The header contains an 8 bit VPI (virtual path
8 index) 42 and a 16 bit VCI (virtual circuit index) 44, which
9 is locally assigned and kept as an index into a look-up
10 table which associates an exit port with this index.

11 Figure 3a shows the GX packet format for the present
12 invention. GX packet 52 comprises a GX header 54, a GX
13 payload 56, and a Frame Check Sequence (FCS) 58, which
14 operates over the entire GX header 54 and GX payload 56

15 Figure 3b shows the details of the GX header of figure
16 3a. GX header 54 includes a START symbol 60. BPDU field 62
17 is a Bridge Protocol Data Unit field, which flags the
18 following data as configuration information used by the
19 spanning tree algorithms such as IEEE 802.1D, or other
20 configuration information which needs to be given higher
21 priority against packet loss during times of network
22 congestion. Packet type field 64 indicates the type of GX
23 payload data, such as ATM cell, Ethernet, FDDI, etc. VLAN
24 field 66 contains VLAN information, PRIORITY field 67
25 indicates priority information, and Application Specific

1 field 68 allows for the optional transmission of information
2 specific to the needs of the data communication system.

3 Figure 4 shows a GX data stream 70 comprising a
4 plurality of GX packets 52, each followed by an inter-packet
5 interval 74 which comprises an END 59 symbol and a variable
6 amount of IDLE 72 time. The END symbol 59 immediately
7 follows the GX FCS 58 of the preceding GX packet 52, and the
8 variable IDLE 72 time allows for the continuous
9 synchronization of the receiver during times when data is
10 not being transmitted.

11 Figure 5 shows examples of different GX payload 56
12 formats, illustrated by the figures 6 through 10.

13 Figure 6 shows the GX payload 56 as a prior art
14 Ethernet packet, comprising Ethernet header 90, and Ethernet
15 payload 92. The Ethernet FCS 26 is not transmitted, since
16 the GX FCS 58 serves this error checking function in GX
17 links.

18 Figure 7 shows a native IP payload comprising an IP
19 frame 98 having an IP header 100, IP source address 102, IP
20 destination address 104, and IP data 106.

21 Figure 8 shows the GX payload 56 comprising an ATM
22 format 108 including a reserved field 110, ATM header 112,
23 and ATM payload 114. The optional reserved field 110 may
24 occur before the ATM header 112 or after the ATM payload
25 114.

1 Figure 9 shows the GX payload 56 comprising an FDDI
2 format 116, including FDDI header 118, and FDDI payload 128.

3 Figure 10 shows the GX payload 56 comprising a Token
4 Ring format 130 including a Token Ring header 132 and Token
5 Ring payload 142.

6 Each of the GX payloads 56 as described in figures
7 6,8,9, and 10 is associated with a particular GX header type
8 field value 64. The objective of field value 64 is to
9 provide knowledge of the type and format of the associated
10 GX payload data. It is clear to one skilled in the art that
11 it is possible to add support for additional payload format
12 types through the assignment of header type 64 beyond those
13 discussed here, and the use of specific GX payload types
14 described herein are not intended to limit the application
15 of field types to only these examples used for illustration.

16 Figure 11 shows the GX packet format for the case where
17 the number of data lanes 152 $n = 8$. Data from the GX packet
18 52 is distributed across 8 data lanes identified by the
19 columns 0 through 7 of 152. Time intervals are identified
20 alternately as odd (o) and even (e) as shown in column 154.
21 K symbols are sent during even cycles, and R symbols sent
22 during odd cycles, as represented by rows 156 and 158,
23 respectively. This alternating sequence continues during
24 inter-packet idle time 167, serving as the variable length
25 inter-packet interval 74. The GX packet 168 comprises GX

1 header 164 and GX payload, starting at row 166, and ending
2 at row 170. Data is instantaneously transmitted from column
3 0 through column 7, row by row, until the final 4 byte FCS
4 58 is sent, as noted by F0, F1, F2, and F3 in row 170, which
5 is the end of the GX packet 52. The T symbol of row 171
6 indicates the end of transmission, and the remainder of the
7 data lanes are filled with odd idle symbol R. The following
8 Ethernet packet 180 of figure 11 comprises GX header 172
9 followed by data, the FCS, and an END symbol in rows 173
10 through 175. If the receiver for the link were busy, the
11 BUSY_IDLE symbols KB and RB would be sent during the inter-
12 packet period 181 from END symbol T on column 7 of line 175
13 followed by BUSY_IDLE symbols on line 176 through 179.

14

15 The algorithm used in transmission is as follows:

- 16 1) Packet data is divided across n data lanes.
- 17 2) Upon startup, each of data lanes 0 through n-1 sends
18 idle symbols alternately coded K during even cycles, and R
19 during odd cycles. The number of these initial startup idle
20 packets is N_{idle} , and their purpose is to gain
21 synchronization of the receiver prior to sending data. If
22 the receiver of the link is busy, and unable to receive
23 additional data, it signals this with the code KB during
24 even cycles, and RB during odd cycles.

3) When a data packet is ready for transmission, a header is sent including the start symbol S on the first data lane, and remainder of the header on the remaining data lanes. The GX header includes information which declares the format of the packet data of the GX payload.

4) Across all data lanes, data is sequentially sent up to, but not including, the last lane-wide data transfer.

5) On the last lane-wide data transfer, the end of packet symbol T is sent on the data lane following the last valid data byte, and the balance of the data lanes is filled with K for even cycles, or R for odd cycles. If the receiver for the link is unable to receive new data, KB is sent during even cycles, and RB is sent during odd cycles.

6) Until the availability of the next complete packet for transmission, the symbol K is sent during even cycles and R is sent during odd cycles across all data lanes. During the time the receiver is unable to receive new data, KB is sent during even cycles, and RB is sent during odd cycles. Additionally, a sequence of at least 32 bytes is sent at least once every $n_{\text{elasticity}}$ data bytes to accommodate clocking differences between systems. The value of $n_{\text{elasticity}}$ is determined by the difference in clock frequency between systems, and is typically 2^{15} bytes or greater.

In the example of figure 11, a link is initialized with the transmission of a minimum number of idle patterns which

1 are transmitted by all data lanes during even cycles as the
2 symbol K and during odd cycles as the symbol R. This idle
3 pattern continues a minimum time to enable symbol
4 synchronization in the receiver. This enables the receiver
5 to acquire lock, and the symbol decoders to operate for an
6 interval on synchronization data, as well as to decode their
7 separate version of odd/even, which the receiver will need
8 for symbol decoding. In this example, a 64 byte Ethernet
9 packet 168 is sent, comprising header 164, and data rows 166
10 through 170, followed by END symbol T and IDLE_ODD R symbols
11 of row 171. Since the last data symbol occurs on the final
12 data lane of 170, data row 171 has the END symbol T sent on
13 data lane 0, and the remainder of the data lanes carry the
14 IDLE_ODD symbol R. The next packet sent is Ethernet packet
15 180, comprising header 172, and data 173 through 175. In
16 this case, the END symbol is sent on the last data lane
17 during the last data transfer 175, which is data lane 7 for
18 the present case of $n=8$.

19 Figure 12 shows the GX packet format for the case where
20 the number of data lanes $n = 4$. As before, the data is
21 presented to data lanes 192, and the frame comprises inter-
22 packet idle pattern 224 which alternates during odd and even
23 cycles of 196, 198, 200, and 202. The GX header is now sent
24 over two cycles starting with the Start symbol in H0 of 204,
25 followed by the balance of the header in second cycle 206.

1 The GX payload 228 is sent from cycles 208 through FCS 216
2 and END symbol T in line 218. Cycles 218 and 220 show a
3 busy receiver, as BUSY_IDLE_EVEN KB and BUSY_IDLE_ODD RB are
4 sent. The following cycle 222 signals an even cycle with
5 the receiver no longer busy.

6 Figure 13 shows the GX packet format for the case where
7 the number of data lanes $n=2$. For this case, the GX header
8 and GX data are distributed across the available data lanes,
9 as shown.

10 Figure 14 shows the serial case of $n=1$, where only one
11 data lane is used. As before, IDLE_ODD and IDLE_EVEN cycles
12 are shown during the inter-packet idle interval 346, and the
13 GX header is sent over 8 cycles as shown in 348, and the GX
14 payload is sequentially sent from 320 through 340, followed
15 by the END symbol T in row 342.

16 It is clear to one skilled in the art that in the
17 previous illustrations for various numbers of byte lanes
18 that the assignment of a particular data byte stream to a
19 particular byte lane is arbitrary, and the sequential
20 assignment shown herein is for illustration only.

21 Examining the header for each of the previous figures
22 11, 12, 13, and 14, figure 3b shows the header format
23 typically used by each of the previous examples. For the
24 present example, where 8 data lanes are used, the header 54
25 comprises the following bit fields:

1 an 8 bit START symbol 60;
2 a single bit BPDU field 62;
3 a 7 bit type field 64 indicating the type of data
4 carried (Ethernet, Token Ring, FDDI, ATM, IP, etc);
5 a 12 bit VLAN field;
6 a 3 bit priority field;
7 a 1 bit reserved field;;
8 a 32 bit application specific field 68 for future expansion.
9 Alternatively, application specific field 68 may contain
10 information that is specific to a particular protocol. The
11 information in the GX header 54 is generally useful to the
12 router or switch in the extraction of early information
13 about the nature and handling of the information to be
14 processed. The type field 64 is used to declare whether the
15 payload data is Ethernet, ATM, or some other type, and the
16 data that follows is used in a context specific to the type
17 declaration 64. In the particular example, a BPDU bit 62
18 can be set which ensures that BPDUs (Bridge Protocol Data
19 Unit) used for Spanning Tree Protocol described in IEEE
20 802.1D properly arrive. Often in networks that are
21 congested because of a reconfiguration event (equipment that
22 has recently failed, or is being added or removed), BPDUs
23 without such priority treatment and which are essential to
24 the reconfiguration of the network are lost, or blocked from

1 transmission. By including this information in the header,
2 they may be granted priority over ordinary data packets.

3 Figure 15 shows a block diagram of the transmit
4 processor 360. Input packets arrive via input interface 364
5 to transmit buffer/controller 370 accompanied by
6 control/data bits 366. Data is arranged in byte order
7 across output ports 0 through n-1, shown for the instance
8 where n=8 having data lanes 0 through 7. In addition to
9 data outputs, the transmit buffer/controller also outputs
10 control information for each data lane. Each data lane is
11 controlled by a single control signal such as 374a for lane
12 0, and accompanied by transmit data such as 372a. During
13 the times that control symbols are emitted, such as START,
14 IDLE, IDLE_BUSY, and END, the control input is asserted
15 true, and the associated 8 bit data indicates which special
16 10B control symbol should be emitted. During the time that
17 the control signal is not asserted, the incoming data is
18 encoded using the prior art 8B/10B coding rule, as will be
19 described later. When IDLE cells are to be output, the
20 encoders across all data lanes are set to alternate between
21 IDLE_EVEN and IDLE_ODD. During the first cycle of a packet
22 transmit, the transmit buffer/controller outputs START on
23 control lane 0. Lanes 1 through 7 carry the balance of the
24 GX Header information. During the last cycle of transmit,
25 all of the lanes up to the final data byte have control

1 signals set to DATA, and data is sent on these data lines.
2 The following lane has its control and data lines set for
3 END, and the remaining lanes have their control and data
4 lines set for IDLE_EVEN, or IDLE_ODD, depending on whether
5 they occur on an even or odd cycle. Encoders 378a through
6 378h accept this control signal accompanied by data. During
7 the times the control input is DATA, this data is encoded
8 from 8B to 10B as described in the 8B/10B coding scheme of
9 U.S. Patent #4,486,739. During the times the control input
10 is CONTROL, the data input is interpreted to produce control
11 codes such as IDLE_EVEN, IDLE_ODD, START, and END. The 10B
12 encoded data is input to transmit serializers 386a through
13 386h, and are output as serial data organized by data lanes
14 388a through 388h. These outputs are typically fed through
15 an optical or electrical link to a receive processor.
16 Transmit clock source 368 provides a clock 376 for the
17 transmit buffer/controller 370 and optionally the encoders
18 378a through 378h at the data rate in use. For a system
19 sending 10Gb/s of data, TX Data 364 would clock 8 bytes of
20 data at a rate of 156.25Mhz into encoders 378 and serializer
21 386 via parallel clock 376. After encoding the data to a 10
22 bit width 380a, the serializer output stages 386 are clocked
23 at 1.5625Ghz. It is understood to one skilled in the art
24 that the actual clocking speed of the interface as provided
25 by generator 368 may be any frequency, as long as the

1 serializer clock rate and transmit buffer output clocking
2 rate are matched in throughput. Additionally, the encoders
3 may be optionally clocked on input or output, or not at all,
4 and the serializers may be optionally clocked on input or
5 output, although best performance may be seen with clocking
6 at each stage.

7 Figure 16 shows an alternate transmit processor for the
8 case where the number of data lanes $n = 4$, while the TX data
9 input 404 rate remains constant. Transmit buffer 410 may
10 accept 8 byte wide inputs on interface 404 accompanied by
11 control/data bits 406, and perform a 2:1 multiplexing to
12 distribute this input data and control across the 4 data
13 lane wide interface. The operation of other elements of
14 figure 16 is the same as the earlier figure 15, with the
15 exception of clock generator 408, which for a 10Gb/s input
16 rate 404, is now clocking in TX data 404 with a 156.25 input
17 clock 409, and is using a 312.5Mhz input clock for control
18 414 and data 412 on lanes 0 to 3 to the encoders 418 and
19 serializer inputs 426. The serializer 426 outputs are
20 clocked using a 3.125Ghz clocking rate 424. The output of
21 this transmit processor is the 4 data lane data stream of
22 figure 12. As is clear to one skilled in the art, many
23 other such modification may be made to the transmit
24 processor, including changing the number of data lanes to
25 any value of one or more, whereby the level of multiplexing

1 of the transmit buffer/controller 410 is related to the
2 width of the input packet bus divided by the number of data
3 lanes on the output side of the buffer/controller. For the
4 case where $n=1$ and a 64 bit packet input is afforded, the
5 input would be multiplexed 8:1, and only data lane 0 would
6 be implemented, producing the output described in figure 14.
7 For $n=1$, the transmit buffer would clock 8 byte data in at
8 156.25Mhz, and send a single stream to the encoder, clocked
9 into and out of the serializer at 1.25Ghz. The serializer
10 rate 422 would be 12.5Ghz. As indicated earlier, any
11 clocking rate could be used, as long as the throughputs are
12 matched between stages.

13 Figure 17a shows the 8B/10B encoder in further detail.
14 Data input 412a and control input 414a produce output 420a
15 which is 10B encoded depending on the data and control
16 inputs, as shown in figure 17b. For the control input 414
17 set to DATA, line 442 shows 8 bit input data directly
18 encoded into 10 bit output data according to the method of
19 patent #4,486,739. Line 440 shows the start transaction,
20 wherein an 8B START input 412 and the control input 414 set
21 to CTRL outputs a Start symbol selected from the available
22 and unused 10B encodings. Similarly, lines 444, 446, 448,
23 449, and 450 show the control input 414 set to CTRL while
24 END, IDLE_EVEN, IDLE_ODD, IDLE_EVEN_BUSY, and IDLE_ODD_BUSY
25 are input as 8B data, which causes unique 10B symbols to be

1 emitted which are reserved from the available and unused 10B
2 encodings. There are many different such unique and unused
3 10B codings, and the best selections are made on the basis
4 of running disparity and hamming distance from other codes.
5 In the best mode the codings (in 10B descriptive format) are

6 START K27.7

7 END K29.7

8 IDLE_EVEN K28.5

9 IDLE_ODD K23.7

10 IDLE_EVEN_BUSY K28.1.

11 IDLE_ODD_BUSY K28.0.

12

13 When the receive processor for this end of the link
14 has a full receive buffer or is no longer capable of
15 receiving additional data, it signals this to the
16 transmitter on the opposite end of the link by emitting
17 EVEN_IDLE_BUSY or ODD_IDLE_BUSY, as shown in lines 449 and
18 450.

19 Figure 18 shows the receive processor 452. Input data
20 is applied as serial streams organized by data lane, shown
21 as 450a through 450h. This data is applied to the receive
22 deserializers 454a through 454h, which output 10 bit wide
23 streams of data 458a through 458h, organized by data lane.
24 These 10B streams are applied to the 10B/8B decoders 456a
25 through 456h, which decode 8 bit data 464a through 464h and

1 control information 466a through 466h, shown for lane 0 as
2 464a and 466a respectively, and repeated identically for all
3 decoders 456a through 456h. Receive elasticity
4 buffer/controller 468 accepts these inputs and organizes
5 packet data to transfer to output port 474 based on
6 receiving a control input 466a asserted with the data input
7 464a having the value START on lane 0, and header data on
8 the remaining lanes 464b and 466a through 464h through 466h,
9 followed by the data packet, followed by the control signal
10 466 asserted accompanied by data 464 having the value END on
11 the appropriate data lane. Thereafter IDLE signals are
12 received. The controller is able to extract a variety of
13 error conditions, such as START appearing on a data lane
14 other than 0, an improperly formed IDLE_EVEN, IDLE_ODD,
15 IDLE_EVEN_BUSY, IDLE_ODD_BUSY or END transaction, and the
16 like. These error conditions are signaled through error
17 output 476. Each deserializer may locally recover a clock
18 459 for use in clocking the decoder 456. The receive
19 elasticity buffer/controller also serves to isolate the
20 clock domains between the input clocking rate controlled by
21 the receive data speed whereby control 466 and data 464
22 signals enter the controller at rates determined by the link
23 far-end sender and buffered data 474 is clocked out of the
24 receive elasticity buffer by local system clock 472. While
25 the receive clock domain boundary is shown in the elasticity

1 buffer 468, it is clear to one skilled in the art that this
2 boundary could be located anywhere in the receive processor.
3 In the example shown in figure 18, data is clocked into
4 decoders 456a-456h by recovered clock 459, which is locked
5 to the frequency of the sending source using clock recovery
6 methods well known in the art. Data is clocked out of the
7 receive buffer by local clock 462, which is derived from
8 system clock 472.

9 Figure 19 shows the case where the receive processor
10 has the number of data lanes $n = 4$. This operates
11 analogously to the receive processor of figure 18, except
12 that the data rate of data presented to the receive
13 buffer/controller on inputs 494 and 496 is doubled.

14 Figure 20a shows the 10B/8B decoder. 10 bit parallel
15 input 458a is converted to 8 bit output data 464a, and 1 bit
16 control 466a. Figure 20b shows the cases 470, 474, 476,
17 478, 480, and 482 for Start, End, Even_Idle, Odd_Idle,
18 Even_Idle_Busy, and Odd_Idle_Busy respectively. Case 472
19 shows 10B input data decoded into 8B output data, and the
20 control output signal asserted as type DATA.

21 In the manner described, data packets may be furnished
22 to a controller, encoded into data lanes, serialized, and
23 transmitted as clock-encoded data. These serial streams of
24 byte-organized data may remotely deserialized, decoded, and
25 re-assembled back into data packets for handling.

1 Additionally, header information may optionally be pre-
2 pended such that information required by the receiving
3 equipment may be easily retrieved during the interval the
4 data packet is being received. The above description of the
5 high speed transmit and receive processors has been drawn to
6 a particular implementation to afford a complete
7 understanding of the invention as practiced in its best
8 mode. Although the invention has been described through the
9 best mode case, it is clear to one skilled in the art that
10 there are many other ways of implementing the described
11 invention. For example, while the figures describe receive
12 and transmit processors having $n=1$ to $n=8$ data lanes, it is
13 clear that an arbitrary number of channels would operate in
14 the same manner, and that the parallel paths and codings do
15 not need to be based on 8-bit bytes. Further, while the
16 encoding and decoding of symbols is based on the standard
17 method of 8B/10B coding of U.S. Patent #4,486,739, other
18 coding/decoding methods could be used. The transmit
19 buffer/controller is shown as generating control signals
20 which indicate the time that particular codes are to be
21 emitted by the encoders, functions which could be performed
22 by a separate controller. While the clocking rates have
23 been chosen to illustrate a 10Gb/s data rate for exemplar
24 purposes, no such clock speed limitation should be inferred

1 in the present invention, which offers the advantages
2 described for any data rate or clocking speed.